arXiv:0806.4391v1 [cs.LG] 26 Jun 2008

Prediction with Expert Advice in Games with Unbounded One-Step Gains

Vladimir V. V'yugin

Institute for Information Transmission Problems, Russian Academy of Sciences,
Bol'shoi Karetnyi per. 19, Moscow GSP-4, 127994, Russia
e-mail vyugin@iitp.ru

Abstract. The games of prediction with expert advice are considered in this paper. We present some modification of Kalai and Vempala algorithm of following the perturbed leader for the case of unrestrictedly large one-step gains. We show that in general case the cumulative gain of any probabilistic prediction algorithm can be much worse than the gain of some expert of the pool. Nevertheless, we give the lower bound for this cumulative gain in general case and construct a universal algorithm which has the optimal performance; we also prove that in case when one-step gains of experts of the pool have "limited deviations" the performance of our algorithm is close to the performance of the best expert.

1 Introduction

Experts algorithms are used for online prediction or repeated decision making or repeated game playing. Any such algorithm is based on a "pool of experts". At any step t, each expert gives its recommendation. From this, a "master decision" is performed. After that, losses (or rewards) s_t^i are assigned to each expert $i = 1,\ldots,m$ by the environment (or adversary). The master algorithm also receives some loss or reward depending on the master decision. The goal of the master algorithm is to perform almost as well as the best expert in hindsight in the long run.

Prediction with Expert Advice considered in this paper proceeds as follows. We are asked to perform sequential actions at times $t = 1,2,\ldots,T$. At each time step t, we observe results of actions of experts in the form of their gains and losses on steps $< t$. After that, at the beginning of the step t *Learner* makes a decision to follow one of these experts, say Expert i. At the end of step t Learner receives the same gain or loss as Expert i at step t.

We use notations and definitions from [5] and [7]. Let $s_{1:t}^i = s_1^i + \ldots + s_t^i$ be the cumulative loss of Expert i at time t. Given $s_{1:t-1}^i$, $i = 1,\ldots,m$, at time t, a natural idea to solve the expert problem is "to follow the leader", i.e. to select the expert i which performed best in the past. The following simple example from Kalai and Vempala [7] shows that Learner can perform much worse than each expert: let the current losses of two experts on steps $t = 1,\ldots,6$ be

$s^1_{1,2,3,4,5,6} = (0,1,0,1,0,1)$ and $s^2_{1,2,3,4,5,6} = (\frac{1}{2},0,1,0,1,0)$. The "Follow Leader" algorithm always chooses the wrong prediction.

The method of following the perturbed leader was discovered by Hannan [4]. Kalai and Vempala [7] rediscovered this method and published a simple proof of the main result of Hannan. They called the algorithm of this type FPL (Following the Perturbed Leader). Hutter and Poland [5] presented a further developments of the FPL algorithm for countable class of experts, arbitrary weights and adaptive learning rate.

The FPL algorithm outputs prediction of an expert i which minimizes

$$s^i_{1:t-1} - \frac{1}{\epsilon}\xi^i_t,$$

where ξ^i_t, $i = 1,\dots m$, $t = 1,2,\dots$, is a sequence of i.i.d random variables distributed according to the exponential distribution with the density $p(t) = e^{-t}$, and ϵ is *a learning rate*. Kalai and Vempala [7] show that the expected cumulative loss of the FPL algorithm has the upper bound

$$E(s_{1:t}) \le (1+\epsilon)\min_{i=1,\dots,n} s^i_{1:t} + \frac{\log n}{\epsilon},$$

where ϵ is a learning rate, n is the number of experts..

In the papers cited above the loss of each expert i can change at any step t by a bounded quantity, for example, $0 \le s^i_t \le 1$ for all t. Poland and Hutter [6] extended this analysis for games with one-step losses upper bounded by an increasing sequence B_t given in advance, i.e., $s_t \le B_t$ for all t. Allenberg et al. [1] also considered unbounded losses, but with different algorithm than in this paper.

In games considered in this paper the players will incur gains (loss is a negative gain); s^i_t denotes one-step gain of a player i. For practical purposes, the property $0 \le s^i_t \le 1$ seems to be too restrictive.

In Appendix A we consider some applications of results of Sections 2-4 of this paper. We define two financial experts learning the fractional Brownian motion whose one-step gains at any step can not be restricted in advance. This application is at the bottom of our special interest in zero-sum games with unbounded gains in Section 4.

In this paper we present some modification of Kalai and Vempala algorithm for the case of unrestrictedly large one-step gains not bounded in advance. We show that in general case, the cumulative gain of any probabilistic prediction algorithm can be much worse than the gain of some expert of the pool. Nevertheless, we give the lower bound for cumulative gain of any probabilistic algorithm in general case and prove that our universal algorithm has optimal performance; we also prove that in case when one-step gains of experts of the pool have "limited deviations" (in particular, when they are bounded) the performance of our algorithm is close to the performance of the best expert. This result is some improvement of results mentioned above.

2 Learning in games of two experts with unbounded gains

In this section we give some preliminary results presenting the bounds on the performance of the algorithm constructed in Section 3.

We consider a simple game G of prediction with expert advice by following of two experts with unbounded one-step gains. The goal of the master algorithm is to receive a cumulative gain not much worse than the gain of the best expert in hindsight.

At each step t of the game both experts receive the nonnegative one-step gains s_t^1 and s_t^2, and their cumulative gains after step t are equal to $s_{1:t}^1 = s_{1:t-1}^1 + s_t^1$ and $s_{1:t}^2 = s_{1:t-1}^2 + s_t^2$.

For simplicity, we consider a variant when at each step t of the game G only one expert can receive a nonnegative one-step gain s_t, and the total gain of the other expert is unchanged, i.e., $s_{1:t}^1 = s_{1:t-1}^1 + s_t$ and $s_{1:t}^2 = s_{1:t-1}^2$ or $s_{1:t}^2 = s_{1:t-1}^2 + s_t$ and $s_{1:t}^1 = s_{1:t-1}^1$. In the general case the analysis is similar.

We also consider *non-degenerate* experts (games), i.e., such that
$\max\{s_{1:t}^1, s_{1:t}^2\} \to \infty$ as $t \to \infty$.

A probabilistic algorithm of following the leader in the game with two experts is based on a computable function f which given cumulative gains $s_{1:t-1}^1$ and $s_{1:t-1}^2$ of the experts in hindside outputs the probability of following the first expert $P\{I = 1\} = f(s_{1:t-1}^1, s_{1:t-1}^2)$ and the probability of following the second expert $P\{I = 2\} = 1 - P\{I = 1\}$.

The analysis in case when these probabilities depend of the whole history of gains is similar.

Let two experts be given. The *master* algorithm works as follows.

Probabilistic algorithm of following the leader.
FOR $t = 1, \dots T$

Given past cumulative gains of the experts $s_{1:t-1}^1$ and $s_{1:t-1}^2$ choose the expert $i \in \{1, 2\}$ with probability $P\{I = i\}$.

Receive the one-step gains at step t of two experts s_t^1 and s_t^2 and define one step gain $s_t = s_t^i$ of the master algorithm.
ENDFOR

The following theorem says that if a probabilistic algorithm of following the leader has high performance in games with bounded one-step gains then its performance in games with unbounded one-step gains can be much worse than the performance of some experts.

Theorem 1. *Let δ, δ' be arbitrary close and arbitrary small positive real numbers such that $\delta' > \delta$, and let for any two non-degenerate experts with bounded one-step gains s_t^i, $i = 1, 2$, i.e., such that $0 \le s_t^i \le 1$ for all t, a master algorithm has the expected cumulative gain*

$$E(s_{1:t}) \ge (1 - \delta) \max_{i=1,2} s_{1:t}^i \tag{1}$$

for all sufficiently large t. Then there exist two experts with unbounded one-step gains such that the expected cumulative gain of the master algorithm is bounded from above

$$E(s_{1:t}) \le \delta' \max_{i=1,2} s^i_{1:t} \tag{2}$$

for infinitely many t.

Proof. Let a master algorithm be given, and let $P\{I = 1\} = f(s^1, s^2)$ and $P\{I = 2\} = 1 - P\{I = 1\}$ be probabilities to choose the best expert from two experts with cumulative gains s^1, s^2. The proof of the theorem uses the following lemma.

Lemma 1. *Let δ, δ' be positive real numbers such that $\delta' > \delta$ and for any two experts with bounded one-step gains the master algorithm has the expected performance (1) for all sufficiently large t. Then for any two real numbers $\tilde{s}^1$ and $\tilde{s}^2$ a number s^1 exists such that $s^1 \ge \tilde{s}^1$, $s^1 \ge \tilde{s}^2$, and $P\{I = 1\} \ge 1 - \delta'$, where $P\{I = 1\} = f(s^1, \tilde{s}^2)$ (and $P\{I = 2\} = 1 - P\{I = 1\}$).*

Proof. Suppose that for some pair $\tilde{s}^1, \tilde{s}^2$ of real numbers the contrary statement holds. Then we can construct two experts with cumulative gains $s^1_{1:t-1}, s^2_{1:t-1}$, $t = 1, 2, \ldots$, and with step-gains equal 0 or 1 such that (1) is violated.

Define the sequences s^1_t, s^2_t, $t = 1, 2, \ldots t_0$, such that s^1_t, s^2_t are equal to 0 or 1 and such that $s^1_{1:t_0} = \tilde{s}^1$ and $s^2_{1:t_0} = \tilde{s}^2$ for some t_0. After that, define $s^1_t = 1$ and $s^2_t = 0$ for all $t > t_0$. We have $s^1_{1:t-1} > s^2_{1:t-1}$ and $P\{I = 1\} < 1 - \delta'$ for all sufficiently large t. Then for the expected one-step gain of the master algorithm,

$$E(s_t) < 1 \circ P\{I = 1\} + 0 \circ P\{I = 2\} \le 1 - \delta' = s^1_t(1 - \delta')$$

holds for all these t. Since $s^1_{1:t} \to \infty$ as $t \to \infty$, we have $E(s_{1:t}) < (1 - \delta)s^1_{1:t}$ for all sufficiently large t. This is a contradiction with (1). Hence, for some t we have $s^1_{1:t-1} > s^2_{1:t-1}$ and $P\{I = 1\} \ge 1 - \delta'$, where $P\{I = 1\} = f(s^1_{1:t-1}, s^2_{1:t-1})$. $\triangle$

We define two experts with unbounded one-step gains as follows. Define $s^1_0 = 0$ and $s^2_0 = 0$. By Lemma 1 a number $s^1 > 0$ exists such that $P\{I = 1\} \ge 1 - \delta$, where $P\{I = 1\} = f(s^1, 0)$. Define $s^1_1 = s^1$, $s^2_1 = 0$.

Let t be even, and let $s^1_{1:t-1}$ and $s^2_{1:t-1}$ be defined on previous steps. We will use the induction hypothesis: $s^1_{1:t-1} > s^2_{1:t-1}$ and $P\{I = 1\} \ge 1 - \delta$. By definition this induction hypothesis holds $t = 2$. Define one-step gains of experts 1 and 2 at step t: $s^1_t = 0$ and $s^2_t = M_t$, where $M_t = \frac{E(s_{1:t-1})}{\delta' - \delta}$ and $E(s_{1:t-1})$ is the mathematical expectation of the cumulative gain of the master algorithm on steps $< t$.

Let t be odd. By Lemma 1 a number s^1 exists such that $s^1 \ge s^1_{1:t-1}$, $s^1 \ge s^2_{1:t-1}$, and $P\{I = 1\} \ge 1 - \delta$, where $P\{I = 1\} = f(s^1, s^2_{1:t-1})$. Define $s^1_t = s^1 - s^1_{1:t-1}$ and $s^2_t = 0$. Then $s^1_{1:t} = s^1$ and $s^2_{1:t} = s^2_{1:t-1}$. Evidently, the induction hypothesis is valid after step t.

Let us prove that this definition is correct. Let t be even. By the induction hypothesis $s^1_{1:t-1} > s^2_{1:t-1}$ and $P\{I = 1\} \ge 1 - \delta$, where $P\{I = 1\} =$

$f(s^1_{1:t-1}, s^2_{1:t-1})$. Then $P\{I = 2\} < \delta$. By definition $s^1_{1:t} = s^1_{1:t-1}$ and $s^2_{1:t} = s^2_{1:t-1} + M_t$. Then we obtain an upper bound for the expected one-step gain of the master algorithm

$$E(s_t) = s^1_t E\{I = 1\} + s^2_t E\{I = 2\} \le \delta M_t.$$

For expected cumulative gain, we have

$$E(s_{1:t}) \le E(s_{1:t-1}) + \delta M_t \le \delta'(s^2_{1:t-1} + M_t) = \delta' s^2_{1:t}. \tag{3}$$

Inequality (3) holds for all even steps t. $\triangle$

Decreasing the lower bound of the performance of a probabilistic algorithm for games with bounded one-step gain functions we can increase it for games with unbounded gain functions. The limit case is given by the following simple example. Evidently, the expected cumulative gain of the probabilistic algorithm which chooses one of two experts with equal probabilities $\frac{1}{2}$ has the lower bound

$$E(s_{1:t}) = \frac{1}{2} s^1_{1:t} + \frac{1}{2} s^2_{1:t} \ge \frac{1}{2} \max_{i=1,2} s^i_{1:t} \tag{4}$$

for $i = 1, 2$.

The following simple diagonal argument shows that the cumulative gain of any probabilistic algorithm of following the leader can be bigger than this bound for some experts, analogously, it can be smaller for some experts.

Proposition 1. *For any δ such that $0 < \delta < 1$ and for any probabilistic algorithm of following the best expert, two experts exist such that the expected cumulative gain of this algorithm satisfies*

$$E(s_{1:t}) \le \frac{1}{2}(1 + \delta) \max_{i=1,2} s^i_{1:t}, \tag{5}$$

for all sufficiently large t, where $s^1_{1:t}$, $s^2_{1:t}$ are cumulative gains of these experts. Analogously, two experts exist such that

$$E(s_{1:t}) \ge \frac{1}{2}(1 - \delta) \max_{i=1,2} s^i_{1:t} \tag{6}$$

for all sufficiently large t.

Proof. Given a probabilistic algorithm of following the best expert and δ such that $0 < \delta < 1$ define recursively the gains of expert 1 and expert 2 at any step t as follows. Let $s^1_{1:t-1}$ and $s^2_{1:t-1}$ be cumulative gains of these experts incurred at steps $< t$. Let $M_t = E(s_{1:t-1})/\delta$, where $E(s_{1:t-1})$ is the expected cumulative gain of the master algorithm in the past.

If $P\{I = 1\} > \frac{1}{2}$ then define $s^1_t = 0$ and $s^2_t = M_t$, and define $s^1_t = M_t$ and $s^2_t = 0$ otherwise. Then $E(s_t) = s^1_t P\{I = 1\} + s^2_t P\{I = 2\} \le \frac{1}{2} M_t$ and $E(s_{1:t}) = E(s^1_{1:t-1}) + E(s_t) \le \frac{1}{2}(1 + \delta) M_t \le \frac{1}{2}(1 + \delta) \max_{i=1,2} s^i_{1:t}$ for all sufficiently large t.

To prove (6) define $s^1_t = M_t$ and $s^2_t = 0$ if $P\{I = 1\} > \frac{1}{2}$, and define $s^1_t = 0$ and $s^2_t = M_t$ otherwise. The following derivation is analogous to the proof of (5). $\triangle$

3 Asymptotically optimal algorithm of following the perturbed leader

In this section we show that the bounds (1) and (2) obtained in Theorem 1 can be achieved by some probabilistic algorithm. More correctly, for any $\delta > 0$ using the method of following the perturbed leader we construct a universal algorithm such that for any δ such that $0 < \delta < 1$ the lower bound

$$E(s_{1:t}) \geq (1-\delta) \max_{i=1,2} s^i_{1:t}$$

is valid for all sufficiently large t for arbitrary two experts ($i = 1, 2$) with bounded one-step gain functions (and even in more general case), and, at the same time, for some $\delta' > 0$ the bound

$$E(s_{1:t}) \geq \delta' \max_{i=1,2} s^i_{1:t}$$

is valid for all experts with arbitrary unbounded one-step gain functions. Here $E(s_{1:t})$ is the cumulative expected gain of the master algorithm.

Note that in this section the cumulative gain is always nonnegative $s^i_{1:t} \geq 0$ for all t and for $i = 1, 2$. In Section 4 we consider the case when the gains can be negative, i.e., experts can incur losses. Recall that, for simplicity, we suppose that at any step t only one expert can receive a positive one-step gain, i.e., $s^1_t = 0$ or $s^2_t = 0$. We denote $s_t = \max\{s^1_t, s^2_t\}$.

Let $\xi^1_1, \xi^2_1, \xi^1_2, \xi^2_2, \ldots$ be a sequence of i.i.d. random variables distributed according to the exponential law with the density $p(t) = e^{-t}$.

We consider the FPL algorithm with learning rate

$$\epsilon_{t-1} = \frac{1}{\mu \max\{s^1_{1:t-1}, s^2_{1:t-1}\}}, \tag{7}$$

where $t = 1, 2, \ldots$ and μ, where $0 < \mu < 1$, is a parameter of the algorithm.

We suppose without loss of generality that $s^1_0 = s^2_0 = 1$. By definition the sequence $\epsilon_1, \epsilon_2, \ldots$ is non-decreasing.

The FPL algorithm is defined as follows:

FPL algorithm.

FOR $t = 1, \ldots T$

Output prediction of expert $i = i_{max}$ which maximizes

$$s^i_{1:t-1} + \frac{1}{\epsilon_{t-1}} \xi^i_t, \tag{8}$$

where $i = 1, 2$, ϵ_{t-1} is defined by (7).

Receive one-step gains s^i_t for experts $i = 1, 2$, and define one step gain $s^{i_{max}}_t$ of the master algorithm.

ENDFOR

Recall that a game G of two experts is called non-degenerate if $v_t = \max\{s^1_{1:t}, s^2_{1:t}\} \to \infty$ as $t \to \infty$, where $s^i_{1:t}$ is the cumulative gain of the expert $i = 1, 2$ at step t. The number

$$\mathrm{Dev}(G) = \limsup_{t\to\infty} \frac{s_t}{v_t}, \tag{9}$$

where $s_t = \max\{s^1_t, s^2_t\}$, is called *the deviation* of the game G. For any game $\mathrm{Dev}(G) \le 1$ by definition. In any non-degenerate game G with bounded one-step gain function, i.e. such that $0 \le s_t \le A$ for all t (A is a positive real number), $\mathrm{Dev}(G) = 0$.

Theorem 2. *For any μ such that $0 < \mu < 1$ an FPL algorithm can be specified such that for any non-degenerate game of two experts its expected cumulative gain at any step T has the lower bound*

$$l_{1:T} \ge e^{-\frac{2}{\mu}}(1-\mu) \max_{i=1,2} s^i_{1:T}, \tag{10}$$

where $s^i_{1:T}$ is the cumulative gain of the expert $i = 1, 2$.

If $\mathrm{Dev}(G) \le \frac{1}{2}\mu\delta$ *for some $0 < \delta < 1$ then*

$$l_{1:T} \ge (1-\delta)(1-\mu) \max_{i=1,2} s^i_{1:T} \tag{11}$$

holds for all sufficiently large T. [1]

Proof. This theorem will follow from Theorem 3 and Corollary 2 below. In the proof we follow the proof-scheme of [5] and [7]. $\triangle$

The analysis of optimality of the FPL algorithm is based on an intermediate predictor IFPL (Infeasible FPL) with the learning rate

$$\epsilon_t = \frac{1}{\mu v_t}, \tag{12}$$

where $v_t = \max\{s^1_{1:t}, s^2_{1:t}\}$.

IFPL algorithm.

FOR $t = 1, \dots T$

Output prediction of expert $i = i_{max}$ with maximal value of

$$s^i_{1:t} + \frac{1}{\epsilon_t}\xi^i_t,$$

where $i = 1, 2$, ϵ_t is defined by (12), and ξ^1_t, ξ^2_t are independent random variables distributed according to the exponential distribution with the density $p(t) = e^{-t}$.

[1] The optimal value of μ for (10) is $\mu = 0.618$. Then $l_{1:T} \ge 0.015 \max_{i=1,2} s^i_{1:T}$ in (10) and $l_{1:T} \ge 0.382(1-\delta)\max_{i=1,2} s^i_{1:T}$ in (11). Comparing these bounds with (4), we reveal a large gap between bounds (10) and (11). Author does not know if we can increase the lower bound (10) when $\mu \approx \frac{1}{2}$ in (11).

Receive one-step gains s_t^i for experts $i = 1, 2$, and define one step gain $s_t^{i_{max}}$ of the master algorithm.
ENDFOR

The IFPL algorithm predicts under the knowledge of $s_{1:t}^1$ and $s_{1:t}^2$ (ϵ_t is their maximum), which both may not be available at beginning of step t. Using unknown value of ϵ_t (like $s_{1:t}^i$, $i = 1, 2$) is the main peculiarity of our version of IFPL.

To distinguish the gains of the FPL and IFPL algorithms we denote s_t^I a one-step gain of the FPL algorithm at step t and s_t^J is a one-step gain of the IFPL algorithm. The expected one-step gains of the FPL and IFPL algorithms at the step t are denoted $l_t = E_t(s_t^I)$ and $r_t = E_t(s_t^J)$.

Theorem 3. *For any μ, $0 < \mu < 1$, the expected one-step gain l_t of the FPL algorithm with learning rate (7) and the expected one-step gain r_t of the IFPL algorithm with learning rate (12) satisfy the inequalities*

$$l_t \geq e^{-\frac{2}{\mu}} r_t \tag{13}$$

for all t.

If $\mathrm{Dev}(G) \leq \frac{1}{2}\mu\delta$ *for some* $0 < \delta < 1$ *then*

$$l_t \geq (1 - \delta) r_t \tag{14}$$

holds for all sufficiently large t.

Proof. For any $t > 0$, denote $\xi^1 = \xi_t^1$, $\xi^2 = \xi_t^2$ and consider two random variables

$$I = \begin{cases} 1 \text{ if } s_{1:t-1}^1 + \frac{1}{\epsilon_{t-1}}\xi^1 > s_{1:t-1}^2 + \frac{1}{\epsilon_{t-1}}\xi^2 \\ 2 \text{ otherwise} \end{cases}$$

and

$$J = \begin{cases} 1 \text{ if } s_{1:t}^1 + \frac{1}{\epsilon_t}\xi^1 > s_{1:t}^2 + \frac{1}{\epsilon_t}\xi^2 \\ 2 \text{ otherwise} \end{cases}$$

Recall that $v_t = \max\{s_{1:t}^1, s_{1:t}^2\}$ for all t. For any real number r we compare conditional probabilities $P\{I = 1|\xi^2 = r\}$ with $P\{J = 1|\xi^2 = r\}$ and $P\{I = 2|\xi^2 = r\}$ with $P\{J = 2|\xi^2 = r\}$.

In our analysis, the nontrivial cases are $s_{1:t}^2 = s_{1:t-1}^2 + s_t$ and $s_{1:t}^1 = s_{1:t-1}^1$ or $s_{1:t}^1 = s_{1:t-1}^1 + s_t$ and $s_{1:t}^2 = s_{1:t-1}^2$, where $s_t > 0$ (we indicate these cases in (16)-(19) below by $\pm$). In this case the following chain of equalities is valid:

$$P\{I = 1|\xi^2 = r\} = P\{s_{1:t-1}^1 + \frac{1}{\epsilon_{t-1}}\xi^1 > s_{1:t-1}^2 + \frac{1}{\epsilon_{t-1}}r|\xi^2 = r\} =$$

$$P\{\xi^1 > \epsilon_{t-1}(s_{1:t-1}^2 - s_{1:t-1}^1) + r|\xi^2 = r\} =$$

$$P\{\xi^1 > \epsilon_t(s_{1:t-1}^2 - s_{1:t-1}^1) + (\epsilon_{t-1} - \epsilon_t)(s_{1:t-1}^2 - s_{1:t-1}^1) + r|\xi^2 = r\} = \tag{15}$$

$$e^{-(\epsilon_{t-1}-\epsilon_t)(s_{1:t-1}^2 - s_{1:t-1}^1)} P\{\xi^1 > \frac{1}{\mu v_t}(s_{1:t-1}^2 - s_{1:t-1}^1) + r|\xi^2 = r\} = \tag{16}$$

$$e^{-\left(\frac{1}{\mu v_{t-1}}-\frac{1}{\mu v_t}\right)(s^2_{1:t-1}-s^1_{1:t-1})\pm\frac{s_t}{\mu v_t}}$$

$$P\{\xi^1 > \frac{1}{\mu v_t}(s^2_{1:t-1} \pm s_t - s^1_{1:t-1}) + r|\xi^2 = r\} = \quad (17)$$

$$e^{\frac{s_t}{\mu v_t}\left(\gamma_t\frac{s^1_{1:t-1}-s^2_{1:t-1}}{v_{t-1}}\pm 1\right)} P\{\xi^1 > \frac{1}{\mu v_t}(s^2_{1:t} - s^1_{1:t}) + r|\xi^2 = r\} = \quad (18)$$

$$e^{\frac{s_t}{\mu v_t}\left(\gamma_t\frac{s^1_{1:t-1}-s^2_{1:t-1}}{v_{t-1}}\pm 1\right)} P\{J = 1|\xi^2 = r\}. \quad (19)$$

Here we have used twice, in (15)-(16) and in (16)-(17), the equality $P\{\xi > a + b\} = e^{-b}P\{\xi > a\}$ for any random variable ξ distributed according to the exponential law; we also used the equality $v_t - v_{t-1} = \gamma_t s_t$, where $0 \le \gamma_t \le 1$, in (16). The exponent (19) is bounded

$$2 \ge \gamma_t\frac{s^1_{1:t-1} - s^2_{1:t-1}}{v_{t-1}} \pm 1 \ge -2. \quad (20)$$

These bounds follow from the inequalities $s^i_{1:t-1}/v_{t-1} \le 1$ and $s^i_{1:t-1} \ge 0$ for all t and for $i = 1, 2$. We also used the inequality $s_t/v_t \le 1$ for all t. Therefore,

$$e^{\frac{2}{\mu}}P\{J = 1|\xi^2 = r\} \ge P\{I = 1|\xi^2 = r\} \ge e^{-\frac{2}{\mu}}P\{J = 1|\xi^2 = r\}. \quad (21)$$

Since, the the inequality (21) holds for all r, it also holds unconditionally

$$e^{\frac{2}{\mu}}P\{J = 1\} \ge P\{I = 1\} \ge e^{-\frac{2}{\mu}}P\{J = 1\}. \quad (22)$$

Analogously, we obtain

$$e^{\frac{2}{\mu}}P\{J = 2\} \ge P\{I = 2\} \ge e^{-\frac{2}{\mu}}P\{J = 2\} \quad (23)$$

for all $t = 1, 2, \ldots$.

If $\mathrm{Dev}(G) \le \frac{1}{2}\mu\delta$, when for sufficiently large t the exponent (19) is bounded from below by

$$e^{-\frac{2}{\mu}\frac{s_t}{v_t}} \ge e^{-\delta} \ge 1 - \delta.$$

From this (14) follows.

From (22) and (23) we obtain the lower bound (13)

$$l_t = E(s^I_t) = s^1_t P(I = 1) + s^2_t P(I = 2) \ge$$
$$s^1_t e^{-\frac{2}{\mu}}P(J = 1) + s^2_t e^{-\frac{2}{\mu}}P(J = 2) = e^{-\frac{2}{\mu}}E(s^J) = e^{-\frac{2}{\mu}}r_t. \quad (24)$$

$\triangle$

The connection between expected cumulative gain of the IFPL algorithm

$$r_{1:T} = \sum_{t=1}^{T} r_t$$

and expected cumulative gain of the FPL algorithm

$$l_{1:T} = \sum_{t=1}^{T} l_t.$$

is given in the following corollary.

Corollary 1. *For any μ and η, $0 < \mu, \eta < 1$, the expected cumulative gains of the IFPL and FPL algorithms with parameters defined in Theorem 3 satisfy the following inequalities*

$$l_{1:T} \geq e^{-\frac{2}{\mu}} r_{1:T} \tag{25}$$

for all T.

If $\mathrm{Dev}(G) \leq \frac{1}{2}\mu\delta$ *for some* $0 < \delta < 1$ *then*

$$l_{1:T} \geq (1-\delta) r_{1:T}$$

holds for all sufficiently large T.

The second bound also holds for unbounded one-step gain games and so, it is some improvement of results of [7] and [5].

The following theorem, which is an analogue of the result from [7], gives a bound for the IFPL algorithm

Theorem 4. *The expected cumulative gain of the IFPL algorithm with the learning rate (12) is bounded by*

$$r_{1:T} \geq \max_{i=1,2} s^i_{1:T} - \frac{1}{\epsilon_T} \tag{26}$$

for all T.

The proof is along the line of the proof from [5] (which is a refinement of the proof from [7]).

Let in this proof $s_t = (s^1_t, s^2_t)$ be a vector of one step gains and $s_{1:t} = (s^1_{1:t}, s^2_{1:t})$ be a vector of cumulative gains of two experts, also let ξ be a vector whose coordinates are random variables ξ^1_t and ξ^2_t. Define $\epsilon_0 = \infty$ and $\tilde{s}_{1:t} = s_{1:t} + \frac{1}{\epsilon_t}\xi_t$ for $t = 1, 2, \ldots$. Consider the one-step gains $\tilde{s}_t = s_t + \xi_t \left(\frac{1}{\epsilon_t} - \frac{1}{\epsilon_{t-1}}\right)$ for the moment. For any vector s and a unit vector d denote

$$M(s) = \mathrm{argmax}_{d \in D}\{d \circ s\},$$

where $D = \{(0,1)^T, (1,0)^T\}$ is the set of two unit vectors of dimension 2 and $\circ$ is the inner product of two vectors.

We first show that

$$\sum_{t=1}^{T} M(\tilde{s}_{1:t}) \circ \tilde{s}_t \geq M(\tilde{s}_{1:T}) \circ \tilde{s}_{1:T}. \tag{27}$$

For $T = 1$ this is obvious. For the induction step from $T-1$ to T we need to show that

$$M(\tilde{s}_{1:T}) \circ \tilde{s}_T \geq M(\tilde{s}_{1:T}) \circ \tilde{s}_{1:T} - M(\tilde{s}_{1:T-1}) \circ \tilde{s}_{1:T-1}.$$

This follows from $\tilde{s}_{1:T} = \tilde{s}_{1:T-1} + \tilde{s}_T$ and $M(\tilde{s}_{1:T}) \circ \tilde{s}_{1:T-1} \leq M(\tilde{s}_{1:T-1}) \circ \tilde{s}_{1:T-1}$.

We rewrite (27) as follows

$$\sum_{t=1}^{T} M(\tilde{s}_{1:t}) \circ s_t \geq M(\tilde{s}_{1:T}) \circ \tilde{s}_{1:T} - \sum_{t=1}^{T} M(\tilde{s}_{1:t}) \circ \xi_t \left(\frac{1}{\epsilon_t} - \frac{1}{\epsilon_{t-1}} \right). \quad (28)$$

By the definition of M we have

$$M(\tilde{s}_{1:T}) \circ \tilde{s}_{1:T} \geq M\left(s_{1:T} + \frac{1}{\epsilon_T} \right) \circ \left(s_{1:T} + \frac{\xi}{\epsilon_T} \right) = \max_d \{ d \circ s_{1:T} \} + M\left(s_{1:T} + \frac{1}{\epsilon_T} \right) \circ \frac{\xi_T}{\epsilon_T}. \quad (29)$$

The expectation of the last term in (29) is equal to $\frac{1}{\epsilon_T}$. We have also

$$\sum_{t=1}^{T} \left(\frac{1}{\epsilon_t} - \frac{1}{\epsilon_{t-1}} \right) M(\tilde{s}_{1:t}) \circ \xi_t \leq \sum_{t=1}^{T} \left(\frac{1}{\epsilon_t} - \frac{1}{\epsilon_{t-1}} \right) M(\xi_t) \circ \xi_t. \quad (30)$$

We have $P\{\max \xi > y\} \leq P\{\xi^1 > y\} + P\{\xi^2 > y\} = 2e^{-y}$. Since

$$E(M(\xi_t) \circ \xi_t) = E(\max\{\xi^1, \xi^2\}) \leq \int_0^\infty 2e^{-y} dy = 2,$$

the expectation of (30) has upper bound $\frac{2}{\epsilon_T}$. Combining the bounds (28)-(30) we obtain (26). $\triangle$.

Corollary 2. *Let μ, $0 < \mu < 1$, be given. If the game of two experts is non-degenerative then the expected cumulative gain of the IFPL algorithm is bounded by*

$$r_{1:T} \geq \max_{i=1,2} s^i_{1:T}(1-\mu).$$

4 Zero sum games

We consider a simplest example of the game of prediction with expert advice with arbitrary positive and negative one-step gains and losses. We apply these results in Appendix A.

We consider a game G of two experts with zero sum, i.e., $s_t^1 = -s_t^2$ at each step t of the game. If a one-step gain is negative it is called a loss. There are no restrictions on the absolute values of s_t^1. Define *a volume* of the game at step t

$$V_t = \sum_{j=1}^{t} |s_j^1|.$$

A game with zero sum is called *non-degenerate* if $\lim\limits_{t\to\infty} V_t = \infty$. Analogously to (9) we consider the deviation of the game G with zero sum

$$\mathrm{Dev}(G) = \limsup_{t\to\infty} \frac{s_t}{V_t},$$

where $s_t = |s_t^1|$ and V_t is the volume of the game.

Evidently, the expected cumulative gain of the algorithm which chooses one of two experts with probability $\frac{1}{2}$ equals zero.

The following proposition is an analogue of Proposition 1.

Proposition 2. *For any probabilistic algorithm of following the best expert, two experts exist such that the expected cumulative gain of this algorithm $E(s_{1:t}) \le 0$ and two experts exist such that $E(s_{1:t}) \ge 0$ for all t.*

Proof. If $P\{I = 1\} > \frac{1}{2}$ define $s_t^1 = 1$, $s_t^2 = -1$ and define $s_t^1 = -1$, $s_t^2 = 0$ otherwise. The following estimates are analogous to that given in the proof of Proposition 1. $\triangle$

The following theorem which is an analogue of the Theorem 1 for games with zero sum shows that if a probabilistic algorithm of the following the leader has high performance in games with bounded one-step gains then its expected cumulative gain in some games with unbounded one-step expert gains can be arbitrary negative.

Theorem 5. *Let L_t be any sequence of positive real numbers, $t = 1, 2, \ldots$. Let δ, δ' be arbitrary close and arbitrary small positive real numbers such that $\delta > \delta'$, and let for any two experts with bounded one-step gains s_t, i.e. such that $0 \le s_t \le 1$ for all t, the expected cumulative gain of the master algorithm has the lower bound*

$$E(s_{1:t}) \ge (1-\delta)|s_{1:t}^1| \tag{31}$$

for all sufficiently large t. Then there exist two non-degenerate experts with unbounded one-step gains such that the expected performance of the master algorithm is bounded from above

$$E(s_{1:t}) \le 2\delta'|s_{1:t}^1| - (1-2\delta')V_t \tag{32}$$

and such that $V_t \ge L_t$ for infinitely many t, where V_t is the volume of the game.

Proof. The proof is similar to the proof of Theorem 1. It uses a modified version of Lemma 1 which is also valid for negative gains with some evident modifications. [2]

Let a master algorithm be given. We define two experts with unbounded one-step gains as follows. Define $s_1^1 = s_1^2 = 0$. By modified version of Lemma 1 a number s^1 exists such that $s^1 > 0$ and $P\{I = 1\} \geq 1 - \delta$, where $P\{I = 1\} = f(s^1, -s^1)$.

Let t be even, and let $s_{1:t-1}^1$ and $s_{1:t-1}^2 = -s_{1:t-1}^1$ be defined on previous steps. We will use the induction hypothesis: $s_{1:t-1}^1 > 0$ and $P\{I = 1\} \geq 1 - \delta$, where $P\{I = 1\} = f(s_{1:t-1}^1, s_{1:t-1}^2)$.

Define one-step gains of experts 1 and 2: $s_t^1 = -M_t$ and $s_t^2 = M_t$, where

$$M_t = \max\left\{\frac{|E(s_{1:t-1})|}{2(\delta - \delta')}, L_t, \frac{V_{t-1}}{\delta}\right\}.$$

Let t be odd. By modified version of Lemma 1 a number s^1 exists such that $s^1 > |s_{1:t-1}^1|$ and $P\{I = 1\} \geq 1 - \delta$, where $P\{I = 1\} = f(s^1, -s^1)$. Define $s_t^1 = s^1 - s_{1:t-1}^1$, then $s_{1:t}^1 = s^1$, and $s_t^2 = -s_t^1$. Evidently, the induction hypothesis is valid after odd step t.

Let us prove that this construction is correct. Let t be even. Then by the induction hypothesis $s_{1:t-1}^1 > 0$ and $P\{I = 1\} \geq 1 - \delta$ (and $P\{I = 2\} < \delta$). By definition $s_{1:t}^1 = s_{1:t-1}^1 - M_t$ and $s_{1:t}^2 = s_{1:t-1}^2 + M_t$. The expected one-step gain of the master algorithm is bounded

$$E(s_t) = s_t^1 P\{I = 1\} + s_t^2 P\{I = 2\} \leq -(1 - \delta)M_t + \delta M_t = -(1 - 2\delta)M_t.$$

By definition $L_t \leq M_t \leq V_t = V_{t-1} + M_t \leq (1 + \delta)M_t$. Then

$$E(s_{1:t}) \leq E(s_{1:t-1}) - (1 - 2\delta)M_t \leq -(1 - 2\delta')M_t = \\ -M_t + 2\delta' M_t \leq -(1 - \delta)V_t + 2\delta'|s_{1:t}^1|$$

for all even steps t. $\triangle$

We consider the *non-degenerate games*, i.e., such that V_t is unbounded.

To obtain the lower bounds we reduce our zero sum game to a game with non-negative one-step gains. Define one-step gain of new experts $\tilde{s}_t^i = s_t^i + |s_t^1|$ for $i = 1, 2$. Then $\tilde{s}_t^i \geq 0$ for all t and $\tilde{s}_t^1 = 0$ or $\tilde{s}_t^2 = 0$ for all t. By definition $\tilde{s}_{1:t}^i = s_{1:t}^i + V_t$ for $i = 1, 2$, where V_t is the volume of the initial game. Evidently, the FPL and IFPL algorithms defined in Section 3 make the same choices for experts of both type.

The expected one-step gains of the master algorithm for for experts of both type satisfy $\tilde{l}_t = s_t^1 P\{I = 1\} + s_t^2 P\{I = 1\} + |s_t|$. This implies the equality $\tilde{l}_{1:t} = l_{1:t} + V_t$ for expected cumulative gains. The analogous equalities hold for $\tilde{r}_t$, $\tilde{r}_{1:t}$ and r_t, $r_{1:t}$.

The following theorem is a corollary of Theorem 2.

[2] A modified version of Lemma 1 looks as follows: Let δ, δ' be positive real numbers such that $\delta' > \delta$, and let for any two experts with bounded one-step gains (31) holds for all sufficiently large t. Then for any number $\tilde{s}^1$ a number $s^1 > 0$ exists such that $s^1 \geq \tilde{s}^1$ and $P\{I = 1\} \geq 1 - \delta'$, where $P\{I = 1\} = f(s^1, -s^1)$.

Theorem 6. *For any μ such that $0 < \mu < 1$ an FPL algorithm can be specified such that for any non-degenerate game of two experts its expected cumulative gain at any step T has the lower bound*

$$l_{1:T} \geq e^{-\frac{2}{\mu}}(1-\mu)|s^1_{1:T}| - V_T(1 - e^{-\frac{2}{\mu}}(1-\mu)), \tag{33}$$

where $s^1_{1:T}$ is the cumulative gain of the the first expert and V_T is the volume of the game at step T.

If $\mathrm{Dev}(G) \leq \frac{1}{2}\mu\delta$ *for some* $0 < \delta < 1$ *then*

$$l_{1:T} \geq (1-\delta)(1-\mu)|s^1_{1:T}| - (\delta+\mu)V_T \tag{34}$$

holds for all sufficiently large T.

Proof. This theorem follows from Theorem 2 and relations between one-step gains $\tilde{s}^i_t$ and s^i_t, $i = 1, 2$, of two type of experts. $\triangle$

Remark. In case when $\mathrm{Dev}(G) \leq \frac{1}{4}\mu\delta$, the bound (34) can be improved for some t if we replace in Section 3 the learning rate (7) on

$$\epsilon_{t-1} = \frac{1}{\mu \max\limits_{j<t} |s^1_{1:j}|}.$$

Then the inequality (34) can be obtained directly (without using the modified experts) from inequalities (22) and (23). We can prove that for any T such that $T = arg \max\limits_{t \leq T} |s^1_{1:t}|$

$$l_{1:T} \geq (1-\delta)(1-\mu)|s^1_{1:T}|.$$

A Learning the fractional Brownian motion

In this section we present some example of the zero sum game studied in Section 4. Rogers [8], Delbaen and Schachermayer [2], and Cheredito [3] have constructed arbitrage strategies for a financial market that consists of money market account and a stock whose price follows a fractional Brownian motion (for continuous time) with drift or an exponential fractional Brownian motion with drift. Vovk [9] has reformulated these strategies for discrete time.

Let $S_0, S_1, \ldots, S_n, \ldots$ be a sequences of prices of some financial instruments such as stocks or bonds. We consider the following "financial" game between an investor and the market. The investor can use the long and short selling.
FOR $t = 1, 2, \ldots T-1$

At the beginning of trading period the investor's cumulative income (or loss) earned from the beginning of the game is $s_{1:t-1} = \sum_{i=1}^{t-1} s_i$.

At the beginning of trading period, observing his past incomes and losses the investor determines the number C_t of shares of the stock needed to realize his strategy.

At the end of trading period the market discloses the price S_{t+1} of the stock, and the investor incur his current income or loss at the period t [3]

$$s_t = C_t(S_{t+1} - S_t).$$

ENDFOR

Denote $\Delta S_t = S_{t+1} - S_t$. We have the following equality

$$(S_T - S_0)^2 = (\sum_{t=0}^{T-1} \Delta S_t)^2 = \sum_{t=0}^{T-1} 2(S_t - S_0)\Delta S_t + \sum_{t=0}^{T-1} (\Delta S_t)^2. \tag{35}$$

The equality (35) leads to the two strategies which are represented by two experts: At the beginning of step t Expert 1 holds the number of shares

$$C_t^1 = 2C(S_t - S_0), \tag{36}$$

Expert 2 holds the number of shares

$$C_t^2 = -2C(S_t - S_0), \tag{37}$$

where C is an arbitrary positive constant.

These strategies at step t earn the incomes $s_t^1 = 2C(S_t - S_0)\Delta S_t$ and $s_t^2 = -s_t^1$. The strategy (36) earns in T steps of the game the income $s_{1:T}^1 = 2C((S_T - S_0)^2 - \sum_{t=1}^{T-1} (\Delta S_t)^2)$. The strategy (37) earns in T steps the income $s_{1:T}^2 = -s_{1:T}^1$.

The number of shares $C_t^1 = 2C(S_{t-1} - S_0)$ in the strategy (36) or number of shares $C_t^2 = -2C(S_{t-1} - S_0)$ in the strategy (37) can be positive or negative. Expert 1 uses the hypothesis that the Hurst exponent of the price of stock is $> \frac{1}{2}$ (a smoother trend). Expert 2 uses the hypothesis that the Hurst exponent is $< \frac{1}{2}$ (volatility is high).

It is reasonable to derandomize the FPL algorithm for this financial game. For that, the investor must follow both experts strategies simultaneously holding $P\{I = 1\}C_t^1 + P\{I = 2\}C_t^2$ shares of a stock at any step t. In this case Theorem 6 holds, where the expected gain at step t is replaced on a pure gain $s_t = P\{I = 1\}C_t^1\Delta S_t + P\{I = 2\}C_t^2\Delta S_t$. [4]

References

1. Chamy Allenberg, Peter Auer, Laszlo Gyorfi and Gyorgy Ottucsak: Hannan Consistency in On-Line Learning in Case of Unbounded Losses Under Partial Monitoring. LNCS, Volume 4264, 229-243, Springer-Verlag Berlin Heidelberg 2006.
2. Delbaen F., Schachermayer W.: A general version of the fundamental theorem of asset pricing. Mathematische Annalen, 300 (1994), 463-520.

[3] We suppose that this price is also valid at the beginning of the period $t + 1$.

[4] Analogously we can derandomize all probabilistic games of this paper if we allow for Learner to receive a given fraction of the gain of an expert.

3. Cheredito P.: Arbitrage in fractional Brownian motion, Finance and Statistics, 7 (4) (2003), 533-553.
4. Hannan J.: Approximation to Bayes risk in repeated plays. In M. Dresher, A.W. Tucker, and P. Wolfe, editors, Contributions to the Theory of Games 3, 97-139, Princeton University Press, 1957.
5. Hutter M., Poland J.: Prediction with expert advice by following the perturbed leader for general weights, (S.Ben-Dawid, J.Case, A.Maruoka (Eds.)): ALT 2004 LNAI 3244, 279-293. Springer-Verlag Berlin Heidelberg 2004.
6. Poland J., Hutter M.: Defensive universal learning with experts. for general weight. (S.Jain, H.U.Simon and E.Tomita (Eds.)): ALT 2005 (S.Jain, H.U.Simon and E.Tomita (Eds.)), LNAI 3734, 356-370. Springer-Verlag Berlin Heidelberg 2005.
7. Kalai A., Vempala S.: Efficient algorithms for online decisions. In Proceedings of the 16th Annual Conference on Learning Theory (COLT-2003), LNAI, 506-521, Berlin, 2003, Springer. Extended version in Journal of Computer and System Sciences, 71, 2005, 291-307.
8. Rogers C.: Arbitrage with fractional Brownian motion. Mathematical Finance, 7 (1997), 95-105.
9. Vovk V.: A game-theoretic explanation of the $\sqrt{dt}$ effect, Working paper #5, 2003, http://www.probabilityandfinance.com